

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

URL Mapping Methods and Systems

Inventor(s):

Rico Mariani

Bassam Tabbara

Ariye M. Cohen

Sanjeev K. Rajan

ATTORNEY'S DOCKET NO. MS1-340USC1

EV317722323

RELATED APPLICATIONS

This application is a *continuation* of a U.S. patent application entitled "URL Mapping Methods and Systems", Ser. No. 09/304,133; by inventors Rico Mariani, Bassam Tabbara, Ariye M. Cohen, and Sanjeev K. Rajan, filed Jun. 30, 1997, hereby incorporated herein by reference.

BACKGROUND

URLs or universal resource locators are used to access resources which can be provided over a communications network such as the Internet. URLs are simply formatted strings that identify particular resources. Typically, a user or client will send a URL to a network server that will respond by sending the requested resource, such as a web page, back to the client. URLs that are provided by a client sometimes have a different, simpler form than those that are used by the server and/or web-site rendering engine that provides the requested resource. For example, a client-provided URL might take the following form: `http://<hostname>/<Abs Path>`, where the "hostname" is a network host domain name or address, and "Abs Path" is the location of the resource requested. The location could be a simple directory path to a .htm file, or a list of parameters that the web site rendering engine can act upon to generate HTML-formatted content. Yet, the URL which is used by the web-site rendering engine is often much more complex in form than the client-provided URL. For example, the form used by the rendering engine might take the following form: `<Abs Path>/script/foo.dll/<identifier>`. The URL form that is used by different rendering engines can really take many forms, and contain the required details that are necessary for the rendering engine to find or generate the requested resource.

1 It is desirable for this more complex URL form to be transparent to the client.
2 This way, the client need only see and/or provide the more simple "user-friendly"
3 URL to the server or web-site rendering engine. Additionally, the underlying
4 complex URL may in fact change over time, while it is desirable for the external
5 form of the URL to remain unchanged so as to not affect customers of the service.

6 Against this backdrop, a need has arisen for a solution to the problem of
7 translating or mapping the simpler, user-friendly URLs into the more complex
8 URLs needed by a rendering engine. Attempts in the past have failed to provide
9 truly flexible, convenient, and dynamic solutions. For example, one past solution
10 was a so-called hard-coded approach which simply provided fixed software code
11 with built-in logic to handle the desired mappings. One way of doing this is to
12 provide a series of IF, THEN statements or a CASE statement to handle the
13 mapping. Problems associated with this approach include that the set of mappings
14 which is supported is fixed, and any changes, such as adding new mapping
15 statements, or modifying existing mapping statements, requires the code to be
16 modified or rewritten. This, in turn, results in a large turn around time which is
17 highly undesirable. In addition, the web service must typically be stopped and
18 restarted which is undesirable and unacceptable for large web-sites that see a large
19 amount of traffic. Moreover, building the rule-mapping logic into the code makes
20 it difficult to have any hierarchical organization of the rules, or to support reverse
21 mapping of the rules. So, in short, the solutions which have been proposed and
22 implemented to date do not provide the flexibility, convenience, and dynamic
23 performance which is so desirable, and in fact, necessary in the current operating
24 environment.

25 SUMMARY

1 The invention provides for flexible, convenient, and dynamic URL
2 mapping methods and systems by providing generalized approaches which specify
3 patterns in term of recognizable syntax. The recognizable syntax can be defined
4 by one or more rules. The rules are used and applied to an input URL string which
5 is provided by a client to map the input URL string to an output URL string which
6 is used by a rendering engine to provide a requested resource.

7 Embodiments of the invention provide for a rule cache or internal store in
8 which rules are kept. The rule cache allows for rules to be added, deleted, or
9 modified as desired, without the need to stop and re-start the web service. The
10 rule cache can be populated dynamically at run time. In addition, embodiments of
11 the invention conveniently group rules into rule groups. Individual rule groups
12 can be selectively applied to an input URL string. Furthermore, reverse mapping
13 is supported by embodiments of the invention simply through the use of suitable
14 rule additions or rule group additions in the rule cache, as may be appropriate.

15 16 **BRIEF DESCRIPTION OF THE DRAWINGS**

17 Fig. 1 shows a communication network which is suitable for use with one
18 or more embodiments of the invention.

19 Fig. 2 is a diagram of a computer that is suitable for use in connection with
20 one or more embodiments of the invention.

21 Fig. 3 is a table which defines so-called special characters in accordance
22 with one or more embodiments of the invention.

23 Fig. 4 is flow diagram which describes methodical aspects in accordance
24 with one or more embodiments of the invention.
25

1 Fig. 5 is a table which provides a collection of rule groups in accordance
2 with one embodiment of the invention.

3 Fig. 6 is table which provides a collection of rule groups in accordance with
4 one embodiment of the invention.

5 Fig. 7 is table which provides a collection of rule groups in accordance with
6 one embodiment of the invention.

7 8 **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS**

9 Fig. 1 shows a communication network 10 in which processing can take
10 place in accordance with the embodiments of the invention described just below.
11 An exemplary communication network is the Internet. The discussion which takes
12 place below assumes some familiarity with the Internet, and World Wide Web
13 practices, formats, and protocols. A great number of books are available on these
14 subjects. Stout, Rick, *The World Wide Web: Complete Reference*, McGraw-Hill,
15 1996, is one example.

16 Communication network 10 includes a server 12 and a client 14. The client
17 and server communicate over a communications medium or link such as the
18 Internet 18 or some other network medium.

19 In the described embodiment, server 12 is an Internet information server. A
20 web-site rendering engine 16 runs on server 12. The web-site rendering engine is
21 responsible for rendering or providing a resource that is requested by client 14.
22 The request for the resource includes an input URL string or user-friendly URL
23 that is in a simple form or easily recognizable form, an example of which was
24 given above.
25

1 A mapping engine 22 also executes on server 12 and initially receives the
2 input URL string. In the illustrated example, mapping engine 22 includes a rule
3 cache 24 and a parser 26. Rule cache 24 includes one or more rules which are
4 applied to the input URL string. Parser 26 performs parsing functions on the input
5 URL string. The result of the operations of rule cache 24 and parser 26 is an
6 output URL string which is provided to web-site rendering engine 16. The output
7 URL string is in a form which can be understood by web-site rendering engine 16.
8 Web-site rendering engine 16 responds by performing the appropriate operations
9 to generate and return the requested resource to client 14 via a response which is
10 sent over communication medium 18.

11 Fig. 2 shows a general example of a desktop computer 130 that can be used
12 in accordance with the invention. A computer such as that shown can be used for
13 any of the client computers 14 and server 12.

14 Computer 130 includes one or more processors or processing units 132, a
15 system memory 134, and a bus 136 that couples various system components
16 including the system memory 134 to processors 132. The bus 136 represents one
17 or more of any of several types of bus structures, including a memory bus or
18 memory controller, a peripheral bus, an accelerated graphics port, and a processor
19 or local bus using any of a variety of bus architectures. The system memory 134
20 includes read only memory (ROM) 138 and random access memory (RAM) 140..
21 A basic input/output system (BIOS) 142, containing the basic routines that help to
22 transfer information between elements within computer 130, such as during start-
23 up, is stored in ROM 138.

24 Computer 130 further includes a hard disk drive 144 for reading from and
25 writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and

1 writing to a removable magnetic disk 148, and an optical disk drive 150 for
2 reading from or writing to a removable optical disk 152 such as a CD ROM or
3 other optical media. The hard disk drive 144, magnetic disk drive 146, and optical
4 disk drive 150 are connected to the bus 136 by an SCSI interface 154 or some
5 other appropriate interface. The drives and their associated computer-readable
6 media provide nonvolatile storage of computer-readable instructions, data
7 structures, program modules and other data for computer 130. Although the
8 exemplary environment described herein employs a hard disk, a removable
9 magnetic disk 148 and a removable optical disk 152, it should be appreciated by
10 those skilled in the art that other types of computer-readable media which can
11 store data that is accessible by a computer, such as magnetic cassettes, flash
12 memory cards, digital video disks, random access memories (RAMs), read only
13 memories (ROMs), and the like, may also be used in the exemplary operating
14 environment.

15 A number of program modules may be stored on the hard disk 144,
16 magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including an
17 operating system 158, one or more application programs 160, other program
18 modules 162, and program data 164. A user may enter commands and information
19 into computer 130 through input devices such as a keyboard 166 and a pointing
20 device 168. Other input devices (not shown) may include a microphone, joystick,
21 game pad, satellite dish, scanner, or the like. These and other input devices are
22 connected to the processing unit 132 through an interface 170 that is coupled to
23 the bus 136. A monitor 172 or other type of display device is also connected to the
24 bus 136 via an interface, such as a video adapter 174. In addition to the monitor,
25

1 personal computers typically include other peripheral output devices (not shown)
2 such as speakers and printers.

3 Computer 130 commonly operates in a networked environment using
4 logical connections to one or more remote computers, such as a remote computer
5 176. The remote computer 176 may be another personal computer, a server, a
6 router, a network PC, a peer device or other common network node, and typically
7 includes many or all of the elements described above relative to computer 130,
8 although only a memory storage device 178 has been illustrated in Fig. 2. The
9 logical connections depicted in Fig. 2 include a local area network (LAN) 180 and
10 a wide area network (WAN) 182. Such networking environments are
11 commonplace in offices, enterprise-wide computer networks, intranets, and the
12 Internet.

13 When used in a LAN networking environment, computer 130 is connected
14 to the local network 180 through a network interface or adapter 184. When used
15 in a WAN networking environment, computer 130 typically includes a modem 186
16 or other means for establishing communications over the wide area network 182,
17 such as the Internet. The modem 186, which may be internal or external, is
18 connected to the bus 136 via a serial port interface 156. In a networked
19 environment, program modules depicted relative to the personal computer 130, or
20 portions thereof, may be stored in the remote memory storage device. It will be
21 appreciated that the network connections shown are exemplary and other means of
22 establishing a communications link between the computers may be used.

23 Generally, the data processors of computer 130 are programmed by means
24 of instructions stored at different times in the various computer-readable storage
25 media of the computer. Programs and operating systems are typically distributed,

1 for example, on floppy disks or CD-ROMs. From there, they are installed or
2 loaded into the secondary memory of a computer. At execution, they are loaded at
3 least partially into the computer's primary electronic memory. The invention
4 described herein includes these and other various types of computer-readable
5 storage media when such media contain instructions or programs for implementing
6 the steps described below in conjunction with a microprocessor or other data
7 processor. The invention also includes the computer itself when programmed
8 according to the methods and techniques described below.

9 For purposes of illustration, programs and other executable program
10 components such as the operating system are illustrated herein as discrete blocks,
11 although it is recognized that such programs and components reside at various
12 times in different storage components of the computer, and are executed by the
13 data processor(s) of the computer.

14 Although Fig. 1 show all of these functions being performed within a single
15 server computer, it is likely that actual embodiments will involve several server
16 computers acting together to satisfy requests of large numbers of clients.
17 Furthermore, the various functions described might be distributed among more
18 than one computer. Also, the mapping engine and web-site rendering engine in the
19 described embodiment are designed to work in conjunction with a Microsoft
20 product called "Internet Information Server." This product performs many of the
21 management functions of a web server, while allowing customization through the
22 use of so-called "filters" and "extensions." Mapping engine 22 is implemented as
23 an ISAPI (Internet Server Application Programming Interface) filter for use in
24 conjunction with Microsoft's Internet Information Server. Web-site rendering
25

1 engine 16 is implemented as an “extension,” again for use with Internet
2 Information Server.

3 As a preliminary matter, the following discussion will use the terminology
4 below:

- 5 • “Input String” is a URL or other string as received by the server
6 from a client. This is typically a user-friendly or “friendly” URL
7 that has been designed for easy recognition and/or recall by human
8 users. This string is passed to the mapping engine for translation
9 into an “unfriendly” format that is appropriate for web-site rendering
10 engine 16.
- 11 • “Output String” is the URL or other string that is output by mapping
12 engine 22 in response to the input string. This string is formatted as
13 appropriate for the web-site rendering engine.
- 14 • “Input Expression” is an expression that is compared with input
15 strings to determine an appropriate mapping. In the described
16 embodiment, an input expression is formatted syntactically in a
17 manner that allows specification of both identity and variability
18 among constituent parts of an input string. Thus, the input
19 expression can include literal parts that call for an exact character-
20 by-character match between those parts and corresponding parts of
21 the input string, and variable parts that allowed for inexact matches
22 or no match at all between those parts and corresponding parts of the
23 input string. An input string is said to “match” an input expression
24 when there is a correspondence between the literal and variable parts
25 of the input string and input expression.

1 “Output Expression” is an expression that is paired with an input
2 expression and that is used to create an output string when there is a
3 match between the input string and the input expression. In the
4 described embodiment, the output expression allows parts of the
5 input string to be specified in the output string. Specifically, the
6 parts of the input string corresponding to variable parts of the input
7 expression can be specified as parts of the output string.

8
9 In any particular server, a plurality of input expressions are defined in
10 accordance with a predefined syntax. The predefined syntax makes use of pattern
11 matching rules. In the described embodiment, this syntax utilizes complex pattern
12 matching rules known as regular expressions. A regular expression comprises a
13 character string in which literal characters indicate text that must exist identically
14 in an input URL string. Regular expressions can also include special characters to
15 indicate portions of an input string in which variability is allowed.

16 As an example, assume that it is desired to map the “friendly” URL input
17 string “seattle.sidewalk.com” to “sidewalk.com/script/foo.dll/seattle”. Assume
18 further that similar mappings are to be made for other cities, such as Portland,
19 Cincinnati, etc.

20 Generally, a matching input string will be any string in which some
21 undefined characters precede the string “.sidewalk.com”. Using a simple form of
22 pattern matching rules, this might be expressed as the following input expression:
23 “*.sidewalk.com”. The “*” indicates any combination of characters, while the
24 following literal characters (“.sidewalk.com”) are to be matched character-by-
25 character with the input string.

1 To produce the appropriate output string, an output expression
2 corresponding to the input string, is formulated as a replacement template. In this
3 example, the output expression might be "sidewalk.com/script/foo.dll/*". The "*" in an output string is an identifier and represents whatever characters corresponded
4 to the "*" in the input string. In this example, the corresponding characters would
5 have been those of the string "seattle". Thus, this output expression would
6 generate the output string "sidewalk.com/script/foo.dll/seattle." It should be
7 apparent that the same input/output expression pair would work with any city
8 specified by a user. This example illustrates a fairly simple and easily
9 understandable syntax. However, more powerful syntax can be used and are often
10 desirable.

12 Fig. 3 shows one example of more a complex set of pattern matching rules
13 known as regular expressions that define a plurality of special characters for
14 specifying variability in an input expression. Regular expressions in accordance
15 with this syntax include so-called escape characters whose meanings are shown.
16 The use of these regular expressions in mapping engine 22 provides a great deal of
17 generality and flexibility in specifying input expressions.

18 Fig. 4 shows a flow diagram generally at 100 that describes certain
19 methodical steps in accordance with an embodiment of the invention. At 102 an
20 input URL string is received by server 12 from a client 14 (Fig. 1). Server 12
21 accesses a plurality of input expressions, each of which is associated with an
22 output expression. At step 104 the input URL string is compared with an input
23 expression. This involves searching the input URL string for a particular pattern
24 that is defined by the input expression. If the input URL string matches the input
25 expression at 106, then the procedure branches to a step 108 of generating an

1 output URL string or pattern from the output expression. The output expression
2 might generate the output URL string by simply causing the mapping engine to
3 conduct a string replacement, e.g. replacing "*" with "seattle", or by doing some
4 additional work such as invoking a lookup procedure. The output URL string is
5 then passed to web-site rendering engine 16, which responds by generating an
6 appropriate web page or other resource. Formulation of the output URL string can
7 take place through multiple transformations, in an iterative fashion, of the input
8 URL string. For example, multiple rules (discussed below) could be used
9 iteratively causing, for example, multiple transformations. If, at 106, the input
10 URL string does not match an input expression, execution proceeds to step 110,
11 which determines whether any more input expressions are available for comparing
12 against the input URL string. If there are, execution loops back to step 104 and
13 106, in which the next input expression is identified and compared with the input
14 URL string. The input expressions might be arranged in a particular hierarchical
15 order so that the input expressions are checked in a stepwise fashion. If there are
16 no more input expressions and there has been no match, then the procedure
17 branches to 112 and returns the URL without modification.

18 In the described embodiment of the invention, mapping engine 22 stores a
19 plurality of rules in its rule cache 24. Each rule comprises an input expression and
20 an output expression. The rules are organized in groups. The rules of a particular
21 group are designed for a particular purpose.

22 Figs. 5-7 show three different groupings of rules that are established or
23 defined using the regular expression syntax defined in Fig. 3. Fig. 5 shows a
24 group of rules that can be used for mapping the hostname part of an input URL
25 string. Fig. 6 shows a rule grouping that can be used for mapping the parameters

1 part (e.g. the "Abs path" portion discussed above). Fig. 7 shows a rule grouping
2 that supports a scoping function which is discussed below in more detail. A rule
3 group consists of a set of zero or more rules. Each rule specifies a mapping of an
4 input URL string to an output URL string using regular expression syntax.

5 Referring specifically to Fig. 5, each rule includes a Rule ID. Here, the
6 Rule IDs number one through twelve. The Rule ID is a number that uniquely
7 identifies a rule inside a rule group. Rules from different groups might take on
8 non-unique IDs, but all rules within a certain group have unique IDs. Rules within
9 each group are applied in the order of the Rule ID, which implies the order in
10 which the rules were added to the group.

11 Each rule is also given a Rule Action Type. Three exemplary Rule Action
12 Types are as follows:

13 "R" Repeat Rule

14 "A" Abort Mapping

15 "N" No Action

16 The Repeat Rule action forces a rule to be applied once again if the rule
17 succeeds. The rule will get applied until the rule fails. The Abort Mapping action
18 implies that if the concerned rule is successfully applied, the mapping process
19 should be immediately aborted and a notification is sent to server 12. The No
20 Action action implies that no special action should be taken if the rule concerned
21 succeeds. That is, the processing of the input URL string should simply continue,
22 i.e. continue following the prescription indicated by the current rule group.

23 Each rule also includes an input expression and an associated output
24 expression. An input expression in the described embodiment is a regular
25 expression in accordance with the syntax of Fig. 3. The output expression is a

1 pattern in accordance with the syntax set forth in Fig. 3. The output expression
2 can be a simple replacement string, or a string including special characters. A \n
3 (where n is a digit [0-9]) in the output expression of a rule corresponds to the nth-
4 tagged expression in the input expression. This provides a convenient notation to
5 extract variable strings from the input expression and insert them into the output
6 expression.

7 A rule is satisfied or succeeds if each of the following events takes place:
8 (1) the input URL string matches the input expression of the rule, based either on a
9 simple string comparison, or a more complex regular expression search; and (2) an
10 attempt to generate the output string based on the output expression succeeds. If a
11 rule is successful, the Rule Action associated with the rule and the Group Tag
12 (discussed below) of the concerned rule group determines what happens next. If
13 this is the last rule in the last group, then this completes the mapping process.

14 Fig. 6 shows a rule grouping that can be used to map parameters of an input
15 URL string. It might be desirable in some applications for the defined syntax to
16 allow invocation of some external procedure such as a lookup procedure. In one
17 embodiment, at least a portion of the output expression can be used to invoke a
18 lookup procedure that produces a result. The result is used to produce an output
19 URL string that is used by the web-site rendering engine to provide the requested
20 resource. The lookup procedure can be invoked in any number of ways. One
21 exemplary way is to instantiate a LookUp object having an interface which
22 supports the lookup procedure. By obtaining a pointer to the interface of the
23 LookUp object, the pointer can be used to invoke the lookup procedure. In this
24 example, the LookUp object might be a COM object. COM objects and interfaces
25 are described in more detail in Brockschmidt, Kraig, *Inside OLE*, Second Edition,

1 Microsoft Press, 1995. It should be noted, however, that the use of COM is in no
2 way necessary to practice the invention, as any suitable object identification
3 schema could be used (e.g. specifying the name of a run-time library file and an
4 entry point, use of a non-COM object directory, etc.). In the illustrated rules, a
5 new special character sequence is shown in the output expressions of Rules 10-13
6 as: “\(\progid, N)”, where “N” is a digit. The sequence invokes a LookUp function
7 that takes the Nth tagged expression as its input. The string argument *progid* is
8 used to create an instance of a LookUp Object that implements an interface
9 ISWFriendlyLookup() that has LookUp() as one of its methods. Specifically, the
10 *progid* obtained from the output expression is used to get the CLSID of the
11 LookUp Object that implements the ISWFriendlyLookup interface. The mapping
12 engine then calls CoCreatesInstance() (a “Windows” operating system call) using
13 this CLSID in order to get a pointer to the ISWFriendlyLookup interface of this
14 LookUp Object. Once the interface pointer is obtained, the mapping engine can
15 cache the pointer for future use. The interface pointer is used to invoke the
16 Lookup() method of that interface. The Lookup() method can take as input, any
17 part of the input URL string for the mapping that matched the tagged expressions
18 specified in the input expressions. The output of the Lookup() method is also a
19 string which replaces the tagged expression in the output expression.

20 Fig. 7 shows a rule group that supports regional scoping. Here, instead of
21 setting a parameter such as “city=seattle” in the URL, a LookUp function is used
22 to resolve an incoming “virtual city name” into a ScopeID instead. The tagged
23 expression syntax in the output expression indicates that a lookup needs to be
24 performed. The mapping engine uses the string “scope” to search a LookUp
25 Cache, and if an entry is found, the cached interface pointer is selected for the

1 Object that implements the lookup method. Otherwise, the mapping engine
2 creates an instance of that particular LookUp Object and then calls the LookUp()
3 method. The LookUp() method can access a database, if necessary, and then
4 returns a scope ID that replaces the tagged expression in the output expression.
5 The rule is then considered successful.

6 As mentioned above, in order to provide for flexibility in the application of
7 rules, the rules can be aggregated in sets to form one or more groups. In essence
8 then, mapping engine 22 comprises a plurality of rule groups, each of which can
9 contain one or more rules. Groups allow aggregation of rules that achieve one
10 aspect of the mapping process into a group. For example, say an input URL
11 stream needs to be decoded by replacing '+' with a space and '%xx' with the
12 character corresponding to the hex value of xx. These operations form two
13 separate rules that take care of one aspect of the URL mapping (that of
14 preprocessing the URL) and can be grouped together in one group.

15 Groups have the following attributes: (1) Group ID, (2) Group Tag, and (3)
16 Group Mask ID. These are shown for the rule groups in Figs. 5 and 6.

17 The Group ID is a unique identifier for the group and is used to identify the
18 group when rules are changed, i.e. added, deleted, or modified. The Group Tag
19 specifies a protocol to use when applying the rules in a particular group. For
20 example, a "match-all" tag specifies that all the rules in a group should be applied.
21 That is, if a rule succeeds, the output of that rule becomes the input of its
22 successor. A "match-one" tag specifies that mapping for that rule group should
23 terminate as soon as a rule matches or succeeds. That is, the output for the rule
24 that succeeds becomes the input of the next group (if one exists) to be considered
25

1 for mapping. If a rule does not succeed, the output of the rule is the same as the
2 input for the rule (as if the rule does not exist).

3 The Group Mask ID is a bit mask which is used during the mapping
4 process. Each group is given a bit mask upon its creation. When a client needs to
5 send a string for mapping, the client can specify a bit-mask key that is logically
6 combined, e.g. ANDed, with the bit mask of each group. If the result of the
7 logical combination is TRUE, then the group is included in the mapping process,
8 otherwise it is skipped. This use of the Group Mask ID provides a convenient
9 method of specifying which groups of rules need to get included in the mapping
10 and which do not. An example of how this feature can be exploited is as follows.

11 It is desirable to have mapping which is bi-directional in the sense that
12 where a given set of rules yields a particular output URL string from a given input
13 URL string, there should be another set of rules in which that particular output
14 URL string yields the given (i.e. the same) input URL string. When doing
15 mapping in one direction, the rules for mapping in the other direction should not
16 apply. One way to achieve this is to provide forward-mapping sets of rule groups
17 with one Group Mask ID, and provide reverse-mapping sets of rule groups with
18 another Group Mask ID. For example, assume groups 1, 2, and 3 have the
19 forward-mapping rules, and groups 4, 5, and 6 have the reverse-mapping rules.
20 Group Mask IDs can then be assigned to the respective groups as follows:

21 Group 1 – 0x0000 0001

22 Group 2 – 0x0000 0002

23 Group 3 – 0x0000 0004

24 Group 4 – 0x0000 0010

25 Group 5 – 0x0000 0020

1 Group 6 – 0x0000 0040
2

3 If the client requests for an input string to be mapped and specifies the bit-
4 mask key to be “0x0000 0007”, then a logical ANDing with the Group Mask IDs
5 would result in only Groups 1, 2, and 3 getting included in the mapping. This
6 would, in turn, provide for forward mapping. On the other hand, a bit-mask key of
7 “0x0000 0070” would result in only Groups 4, 5, and 6 being considered in the
8 mapping, and not groups 1, 2, and 3. In this case, reverse mapping would be
9 applied.

10 As another example, consider a situation where rule groups are desired to
11 be removed. If you want to change the rules, the best way to do it is to add the
12 new rules (using bitmasks not currently in use), then once they are all successfully
13 added, switch to using the new bitmasks and stop using the old ones.
14 Subsequently, the old rule groups can be removed. In this way there is a clean
15 move over to the new rule set.

16 In one specific implementation, a mapping system is provided through the
17 use of Microsoft’s Internet Information Server’s extension facilities and
18 Microsoft’s COM methodologies mentioned above. For purposes in assisting in
19 understanding this implementation, components of the mapping system and
20 pertinent interfaces are described.

21 The mapping engine is implemented as an ISAPI filter. Hence, it
22 implements the functions every ISAPI filter needs to implement, e.g.
23 GetFilterVersion() and HttpFilterProc().

24 **GetFilterVersion():** This function is called only once when IIS (Internet
25 Information Server) is started. It is used for exchanging version information

1 between IIS and the mapping engine. The mapping engine also informs IIS of the
2 notifications that it is interested in. Later, when a certain event occurs, IIS will
3 invoke only those filters that have requested notification for that event.

4 **HttpFilterProc():** IIS notifies the mapping engine by calling
5 HttpFilterProc() and passing it the notification types and a pointer to the structure
6 corresponding to that notification. Using this pointer the mapping engine will gain
7 access to the header information such as <hostname> and <abs_path> (the two
8 parts of the URL discussed above).

9 The mapping engine is a COM server. This means that it can provide its
10 services to anyone who can obtain a pointer to its appropriate interface. This
11 results in generic, rule-based mapping capabilities. Some of the interfaces that can
12 be used to populate and use the mapping engine are as follows:

14 **SWFriendly Interface**

15 This interface provides various methods for managing and using the
16 mapping engine's rules. It provides support for adding new groups of rules to the
17 rule cache and new rules to the groups. It can also provide support for loading
18 rules from a file, and for storing rules back to a file so that the rules persist even
19 after IIS is restarted. Other methods that can be supported can include methods
20 for modifying and deleting rules. For example, "remove" operations that are
21 analogous to the "Add" operations discussed just below can readily be added. A
22 few examples of such methods are given below.

23 **AddGroup()** – This method is used by the clients to add a new group to the
24 rule cache. The client specifies a group tag and group mask Id as input
25

1 parameters, and the method returns a group Id for the newly added group in the
2 output parameter. The client can then use this group Id to add rules to this group.

```
3 HRESULT AddGroup
  //-----
4 // Add a new group (or cache) of rules to Friendly
  (
    EGroupTag eGroupTag,
5    // [in] the group Tag (match one or match all)
    DWORD dwGroupMaskID,
6    // [in] the group bit-mask ID. Used to decide whether
    // group is to be included in the mapping process
    DWORD * pdwGroupID
7    // [out] the groupID generated for the new group
  );
8
```

9 **AddRule()** – This method is used to add a rule to a group that is identified
10 by its group Id. The client provides the following: group Id of the group to which
11 the rule is to be added, the action associated with the rule, the input expression and
12 the output expression for the rule.

```
13 HRESULT CSWFriendly::AddRule
  //-----
14 // Add a new rule to a certain Group.
  (
15    DWORD dwGroupID,
    // [in] GroupID identifying the group to which the
    // rule should be added
16    ERuleAction eRuleAction,
    // [in] Action associated with this rule
17    LPCSTR pszSrcString,
    // [in] Input Expression for the rule
18    LPCSTR pszDestString
    // [in] Output Expression for the rule
  );
19
```

20 **ProcessURL()** – This method is called by the client to send an input URL
21 string to the mapping engine. The client specifies a bit-mask that is used to decide
which groups of rules are to be included for mapping.

```
22 HRESULT CSWFriendly::ProcessUrl
  //-----
23 // Takes the incoming Url as input and sends it through
  // the mapping engine. The input Url gets sent through each
  // group of rules.
24 (
    LPCSTR pszUrl,
25    // [in] Incoming Url
    DWORD dwBitMask,
```

```
1 // [in] this parameter decides the set of groups that
2 // get used for mapping
3 LPSTR pszNewUrl
4 // [in, out] Out Url
5 );
```

ISWFriendlyLookup Interface:

6 This interface provides a method for conducting a lookup procedure. Here,
7 a rule can invoke a lookup procedure that gets its input from the input URL string.
8 The output of this procedure is used in the output generated by the rule, e.g. an
9 output URL string. Objects that implement this interface are referred to as
10 Friendly Lookup Objects or FLOs. Each FLO has a CLSID and a ProgID and
11 registers itself in the registry on compilation. A cache is provided and is referred
12 to as the Friendly Lookup cache. The Friendly Lookup cache contains two fields
13 in each entry – a progid (or a string) and a pointer to the ISWFriendlyLookup
14 interface. When a rule needs to invoke the Lookup method of a FLO, it will use
15 the FLO's ProgID to create an instance of it and will then cache the pointer to the
16 interface in its cache for future references. Specifically, if the output expression of
17 a rule has a tagged expression with the following syntax, it means that the rule
18 requires a Lookup method to be invoked:

19
20 \ (Foo, N) – where Foo is any string and N is an integer.

21
22 If the input URL string matches the input expression for a certain rule, and
23 the output expression contains a tagged expression of the type shown above, the
24 mapping engine performs the following steps:

1. It extracts the string in the tagged expression ("Foo" in the above case) and does a search in its lookup cache based on this string.
2. A call to CLSIDFromProgID() is made to obtain the CLSID of the FLO. Using this CLSID, an instance of the FLO is created by calling CoCreateInstance(). The pointer to the interface obtained in this way is cached in the lookup cache along with the ProgID for future lookups.
3. Using the interface pointer, the Lookup method of this interface is invoked. The input for the lookup method is the sub-string from the input URL string that matches the Nth tagged expression in the input URL string of the rule. Syntax might also be provided in the tagged expression to specify more than one string obtained from the input URL string to be passed to Lookup method. This can provide for more powerful lookups.
4. The output or result of the Lookup method replaces the tagged expression in the output expression.
5. If the Lookup fails, the rule is considered unsuccessful and processing continues just like it would after a rule fails.

The ISWFriendlyLookup interface has the following method:

Lookup() – This method takes as input a string (obtained as explained above) and returns a pointer to a string buffer pointer. The callee allocates memory to hold the output of the Lookup() and the caller frees the buffer.

```
HRESULT Lookup
(
    [in] LPCWSTR pwszSource,
    // pointer to a buffer containing the string that
    // the Lookup method needs for doing the lookup
    [in, out] LPWSTR * ppwszDestination
```

```

1      // a pointer to a buffer that will hold a pointer
2      // to the actual string allocated by the callee.
3      // The callee must use IMalloc:Alloc to allocate
4      // the memory.
5      );

```

Various embodiments of the invention described above provide for a flexible and generic solution to the problem of mapping input URL strings to output URL strings. Rules for mapping can now be changed, i.e. added, deleted, or modified dynamically, without the need to access and rewrite code, or shut down the communication network server. And, while the described embodiments have been described in terms of processing input URL strings, it is possible that other inputs can be used as well. For example, various mapping methods and systems can use other information as inputs, such as that gleaned from an http header. Examples of such other information include cookies, user agent, user browser capabilities and the like. Typically, these are provided as strings. Thus, in these methods and systems, there are two or more inputs to the mapping engine. For example, the two methods specified just below can enable the use of two regular expressions (e.g. one for the URL, and one to process on all of the headers) in order to process the rules.

```

18      HRESULT AddRuleWithHeader
19      (
20          [in] DWORD dwGroupID,
21          [in] ERuleAction eRuleAction,
22          [in] LPCSTR pszSrcString,
23          [in, optional] LPCSTR pszSrcHeaderString,
24          [in] LPCSTR pszDestString,
25          [in] LPCSTR pszAccelerator
26      );
27
28      HRESULT ProcessUrlWithHeader
29      (
30          [in] LPCSTR pszUrl,
31          [in, optional] LPCSTR pszHeaders,

```



```
[in] DWORD dwGroupMaskID,  
[in, out] LPSTR pszNewUrl,  
[in] UINT cchSize  
);
```

Other advantages will be apparent to those of skill in the art.

Although the invention has been described in language specific to structural features and/or methodological steps, it is to be understood that the invention defined in the appended claims is not necessarily limited to the specific features or steps described. Rather, the specific features and steps are disclosed as preferred forms of implementing the claimed invention.